

Local Linear Convergence of the ADMM/Douglas–Rachford Algorithms without Strong Convexity and Application to Statistical Imaging*

Timo Aspelmeier[†] C. Charitha[‡] and D. Russell Luke[§]

version March 16, 2016

Abstract

We consider the problem of minimizing the sum of a convex function and a convex function composed with an injective linear mapping. For such problems, subject to a coercivity condition at fixed points of the corresponding Picard iteration, iterates of the alternating directions method of multipliers converge locally linearly to points from which the solution to the original problem can be computed. Our proof strategy uses duality and strong metric subregularity of the Douglas–Rachford fixed point mapping. Our analysis does not require strong convexity and yields error bounds to the set of model solutions. We show in particular that convex piecewise linear-quadratic functions naturally satisfy the requirements of the theory, guaranteeing eventual linear convergence of both the Douglas–Rachford algorithm and the alternating directions method of multipliers for this class of objectives under mild assumptions on the set of fixed points. We demonstrate this result on quantitative image deconvolution and denoising with multiresolution statistical constraints.

Research of T. Aspelmeier, C. Charitha and D. R. Luke was supported in part by the German Research Foundation grant SFB755-A4.

2010 Mathematics Subject Classification: Primary 49J52, 49M20, 90C26; Secondary 15A29, 47H09, 65K05, 65K10, 94A08.

Keywords: Augmented Lagrangian, ADMM, Douglas–Rachford, exact penalization, fixed point theory, image processing, inverse problems, metric regularity, statistical multiscale analysis, piecewise linear-quadratic, linear convergence

*To appear, SIAM Journal on Imaging Sciences

[†]Institute for Mathematical Stochastics and Felix Bernstein Institute for Mathematical Statistics in the Biosciences, Georg-August Universität Göttingen, 37077 Göttingen, Germany

[‡]Institut für Numerische und Angewandte Mathematik, Universität Göttingen, 37083 Göttingen, Germany

[§]Institut für Numerische und Angewandte Mathematik, Universität Göttingen, 37083 Göttingen, Germany

1 Introduction.

The alternating directions method of multipliers (ADMM) has received a great deal of attention recently for large-scale problems involving constraints on the image of the unknowns under some linear mapping. The analysis has focused on either global complexity estimates [34] or sufficient conditions for local linear convergence [14, 24, 48]. The closely related Douglas–Rachford algorithm has also been the focus of recent studies showing global complexity [42, 52] and (local linear) convergence in increasingly inhospitable settings [1, 7–9, 13, 36, 37, 54]. A survey of results on proximal methods in general can be found in [51]. In the convex setting, the convergence studies for both ADMM and Douglas–Rachford share a common thread through the well-known duality between these algorithms [29]. Studies of ADMM frequently invoke strong convexity. Studies of Douglas–Rachford, on the other hand have, until very recently, focused on *feasibility* problems and corresponding notions of regularity of intersections. We combine an analysis of the ADMM algorithm with facts learned from the local convergence of Douglas–Rachford to provide sufficient conditions for local linear convergence of sequences generated by ADMM without strong convexity. While this paper was under review we became aware of two recent studies that also combine the analysis of ADMM and Douglas–Rachford to improve and generalize many local and global results [30, 31]. While our theoretical results are general and abstract, our motivation for the current study comes from the application of statistical multiscale image denoising/deconvolution following [26, 27] for fluorescence microscopic images (see also [2] for a review of fluorescence microscopy techniques and statistical methods for them). We demonstrate the analysis for image denoising/deconvolution of Stimulated Emission Depletion (STED) images [35, 38].

1.1 Notation and definitions

Though many of the arguments presented here work equally well for infinite dimensional Hilbert spaces, to avoid technicalities, *it will be assumed throughout that U and V are Euclidean spaces*. The norm $\|\cdot\|$ denotes the Euclidean norm. We denote the extended reals by $(-\infty, +\infty] := \mathbb{R} \cup \{+\infty\}$ and the nonnegative orthant by $\mathbb{R}_+ := \{x \in \mathbb{R} \mid x \geq 0\}$. The closed unit ball centered at the origin is denoted by \mathbb{B} . In the usual notation for the natural numbers \mathbb{N} we include 0. The mapping $A : U \rightarrow V$ is linear and the functional $J : U \rightarrow (-\infty, +\infty]$ is proper (not everywhere $+\infty$ and nowhere $-\infty$), convex and lower semicontinuous (lsc), as is the functional $H : V \rightarrow (-\infty, +\infty]$. The level set of J corresponding to $\alpha \in \mathbb{R}$ is defined by $\text{lev}_{\leq \alpha} J := \{u \in U : J(u) \leq \alpha\}$. The domain of a function $f : U \rightarrow (-\infty, +\infty]$ is defined by $\text{dom } f = \{u \in U : f(u) < \infty\}$. We use the notation $\Phi : U \rightrightarrows V$ to denote a set-valued mapping from U to V .

A proper function $f : U \rightarrow (-\infty, +\infty]$ is *strongly convex* if there is a constant $\mu > 0$ such that

$$f((1 - \tau)x_0 + \tau x_1) \leq (1 - \tau)f(x_0) + \tau f(x_1) - \frac{1}{2}\mu\tau(1 - \tau)\|x_0 - x_1\|^2, \quad (1.1)$$

for all x_0 and x_1 and $\tau \in (0, 1)$. We will not assume smoothness of functions and so

will require the *subdifferential*. The subdifferential of a function $f : U \rightarrow (-\infty, +\infty]$ at a point $\bar{x} \in \text{dom } f$ is defined by

$$\partial f(\bar{x}) := \{v \in U \mid \langle v, x - \bar{x} \rangle \leq f(x) - f(\bar{x}), \text{ for all } x \in U\}. \quad (1.2)$$

When $\bar{x} \notin \text{dom } f$ the subdifferential is defined to be empty. Elements from the subdifferential are called *subgradients*. The subdifferential of a proper, lsc convex function is a *maximally monotone* set-valued mapping [56, Theorem 12.17]. The *Fenchel conjugate* of a function f is denoted by f^* and defined by

$$f^*(y) := \sup_{x \in U} \{\langle y, x \rangle - f(x)\}.$$

A mapping $\Phi : V \rightrightarrows V$ is said to be β -*inverse strongly monotone* [56, Corollary 12.55] if for all $x, x' \in V$

$$\langle v - v', x - x' \rangle \geq \beta \|v - v'\|^2, \quad \text{whenever } v \in \Phi(x), v' \in \Phi(x'). \quad (1.3)$$

The mapping Φ is said to be *polyhedral* (or piecewise polyhedral [56]) if its graph is the union of finitely many sets that are polyhedral convex in $U \times V$ [20]. We denote the *resolvent* of Φ by $\mathcal{J}_\Phi := (\text{Id} + \Phi)^{-1}$ where Id denotes the identity mapping and the inverse is defined as

$$\Phi^{-1}(y) := \{x \in U \mid y \in \Phi(x)\}. \quad (1.4)$$

The corresponding *reflector* is defined by $R_{\eta\Phi} := 2\mathcal{J}_{\eta\Phi} - \text{Id}$.

Notions of *continuity* of set-valued mappings have been thoroughly developed over the last 40 years. Readers are referred to the monographs [4, 20, 56] for basic results. A mapping $\Phi : U \rightrightarrows V$ is said to be *Lipschitz continuous* if it is closed valued and for all $u, u' \in U$ there exists a $\tau \geq 0$ such that

$$\Phi(u') \subset \Phi(u) + \tau\|u' - u\|\mathbb{B}. \quad (1.5)$$

Lipschitz continuity is, however, too strong a notion for set-valued mappings. A key property of set-valued mappings that we will rely on is *metric subregularity*, which can be understood as the property corresponding to a Lipschitz-like continuity of the inverse mapping relative to a specific point. As the name suggests, it is a weaker property than *metric regularity* which, in the case of an $n \times m$ matrix for instance, is equivalent to surjectivity. Our definition follows the characterization of this property given in [20, Exercise 3H.4].

Definition 1.1 ((strong) metric subregularity).

- (i) *The mapping $\Phi : U \rightrightarrows V$ is called metrically subregular at \bar{x} for \bar{y} relative to $W \subset U$ if $(\bar{x}, \bar{y}) \in \text{gph } \Phi$ and there is a constant $c > 0$ and neighborhoods \mathcal{O} of \bar{x} such that*

$$\text{dist}(x, \Phi^{-1}(\bar{y}) \cap W) \leq c \text{dist}(\bar{y}, \Phi(x)) \quad \forall x \in \mathcal{O} \cap W. \quad (1.6)$$

- (ii) The mapping Φ is called *strongly metrically subregular* at \bar{x} for \bar{y} relative to $W \subset U$ if $(\bar{x}, \bar{y}) \in \text{gph } \Phi$ and there is a constant $c > 0$ and neighborhoods \mathcal{O} of \bar{x} such that

$$\|x - \bar{x}\| \leq c \text{dist}(\bar{y}, \Phi(x)) \quad \forall x \in \mathcal{O} \cap W. \quad (1.7)$$

The constant c measures the stability under perturbations of inclusion $\bar{y} \in \Phi(\bar{x})$.

An important instance where metric subregularity comes for free is for polyhedral mappings.

Proposition 1.2 (polyhedrality implies strong metric subregularity). *Let $W \subset V$ be an affine subspace and $T : W \rightrightarrows W$. If T is polyhedral and $\text{Fix } T \cap W$ is an isolated point, $\{\bar{x}\}$, then $\text{Id} - T : W \rightrightarrows (W - \bar{x})$ is strongly metrically subregular, hence metrically subregular, at \bar{x} for 0 relative to W .*

Proof. If T is polyhedral, so is $\Phi^{-1} := (\text{Id} - T)^{-1}$. Now by [20, Propositions 3I.1 and 3I.2], since Φ^{-1} is polyhedral and \bar{x} is an isolated point of $\Phi^{-1}(0) \cap W$, then $\Phi = \text{Id} - T$ is *strongly metrically subregular* at \bar{x} for 0 with constant c on the neighborhood \mathcal{O} of \bar{x} restricted to W (1.7). \square

One prevalent source of polyhedral mappings is the subdifferential of piecewise linear-quadratic functions (see Proposition 2.6 below).

Definition 1.3 (piecewise linear-quadratic functions). *A function $f : \mathbb{R}^n \rightarrow [-\infty, +\infty]$ is called *piecewise linear-quadratic* if $\text{dom } f$ can be represented as the union of finitely many polyhedral sets, relative to each of which $f(x)$ is given by an expression of the form $\frac{1}{2}\langle x, Ax \rangle + \langle a, x \rangle + \alpha$ for some scalar $\alpha \in \mathbb{R}$ vector $a \in \mathbb{R}^n$, and symmetric matrix $A \in \mathbb{R}^{n \times n}$.*

A notion related to metric regularity is that of *weak-sharp solutions*. This will be used in the development of error bounds (Theorem 3.4).

Definition 1.4 (weak sharp minimum [16]). *The solution set $\text{argmin } \{f(x) \mid x \in \Omega\}$ for a nonempty closed convex set Ω , is weakly sharp if, for $\bar{p} = \inf_{\Omega} f$, there exists a positive number α (sharpness constant) such that*

$$f(x) \geq \bar{p} + \alpha \text{dist}(x, S_f) \quad \forall x \in \Omega.$$

Similarly, the solution set S_f is weakly sharp of order $\nu > 0$ if there exists a positive number α (sharpness constant) such that, for each $x \in \Omega$,

$$f(x) \geq \bar{p} + \alpha \text{dist}(x, S_f)^{\nu} \quad \forall x \in \Omega.$$

1.2 Preparatory abstract results

To conclude this section we present general results about types of (firmly) nonexpansive operators that clarify the underlying mechanisms yielding linear convergence of many algorithms. The operative definitions are given here.

Definition 1.5 $((S, \epsilon)$ -firmly-)nonexpansive mappings). Let D and S be nonempty subsets of U and let T be a (multi-valued) mapping from D to U .

(i) T is called (S, ϵ) -nonexpansive on D if

$$\begin{aligned} \|x_+ - \bar{x}_+\| &\leq \sqrt{1 + \epsilon} \|x - \bar{x}\|, \\ \forall x \in D, \forall \bar{x} \in S, \forall x_+ \in Tx, \forall \bar{x}_+ \in T\bar{x}. \end{aligned} \quad (1.8)$$

If (1.8) holds with $\epsilon = 0$ then we say that T is S -nonexpansive on D .

(ii) T is called (S, ϵ) -firmly nonexpansive on D if

$$\begin{aligned} \|x_+ - \bar{x}_+\|^2 + \|(x - x_+) - (\bar{x} - \bar{x}_+)\|^2 &\leq (1 + \epsilon) \|x - \bar{x}\|^2, \\ \forall x \in D, \forall \bar{x} \in S, \forall x_+ \in Tx, \forall \bar{x}_+ \in T\bar{x}. \end{aligned} \quad (1.9)$$

If (1.9) holds with $\epsilon = 0$ then we say that T is S -firmly nonexpansive on D . If, in addition, $S = \text{Fix } T$, then T is said to be quasi-firmly nonexpansive.

Theorem 1.6 (abstract linear convergence result). Let $W \subset V$ be an affine subspace and $T : W \rightrightarrows W$ be quasi-firmly nonexpansive on W . Let $\text{Fix } T \cap W$ be an isolated point, $\{\bar{x}\}$. If $\text{Id} - T : W \rightrightarrows (W - \bar{x})$ is metrically subregular at \bar{x} for 0, then there is a neighborhood \mathcal{O} of \bar{x} such that

$$\text{dist}(x_+, \text{Fix } T) \leq \sqrt{1 - \kappa} \text{dist}(x, \text{Fix } T), \quad \forall x_+ \in Tx, \forall x \in \mathcal{O} \cap W, \quad (1.10)$$

where $0 < \kappa = c^{-2}$ for c a constant of metric subregularity of $\text{Id} - T$ at \bar{x} for the neighborhood \mathcal{O} . Consequently, the fixed point iteration $x^{k+1} = Tx^k$ converges linearly to $\text{Fix } T$ with rate $\sqrt{1 - \kappa}$ for all $x^0 \in \mathcal{O} \cap W$.

Proof. Define $\Phi := (\text{Id} - T)$ and note that $\{\bar{x}\} = (\text{Id} - T)^{-1}(0) \iff \{\bar{x}\} = \text{Fix } T$, hence

$$\text{dist}(x, (\text{Id} - T)^{-1}(0)) = \text{dist}(x, \text{Fix } T) = \|x - \bar{x}\|.$$

Suppose that Φ is metrically subregular at $\text{Fix } T$ for 0. Then by Definition 1.1(i) we have, for all $x \in \mathcal{O} \cap W$ and for all $x^+ \in T(x)$,

$$\text{dist}(x, (\text{Id} - T)^{-1}(0)) = \|x - \bar{x}\| \leq c \text{dist}(0, (x - Tx)) \leq c\|x - x^+\|, \quad (1.11)$$

which is the coercivity condition of [36, Eq.(3.1), Lemma 3.1]. By assumption, T is $(\text{Fix } T, 0)$ -firmly nonexpansive (i.e., quasi-firmly nonexpansive) on W (Definition 1.5 (ii)). The result then follows from [36, Lemma 3.1] with rate $\sqrt{1 - \kappa}$ for $\kappa = c^{-2}$. \square

Remark 1.7 (on κ). The constant κ in the above theorem can always be chosen to be less than or equal to 1. To see this, note that for any metrically subregular mapping Φ , there is a constant $c \geq 1$ and hence a $\kappa \leq 1$ so that the rate constant given in Theorem 1.6 will always hold whenever the fixed point is a (relatively) isolated point.

Example 1.8 (a simple example). Consider two lines, A and B , in \mathbb{R}^2 intersecting orthogonally at the origin and let T be the Douglas–Rachford operator for the projections onto each line. In this example $T = \frac{1}{2}(R_A R_B + \text{Id})$ where $R_A := 2P_A - \text{Id}$ for the projection onto the line A denoted by P_A , and likewise for R_B . In the context of what follows, P_A is the resolvent of the subdifferential of the indicator function of the line A and likewise for P_B . It is elementary to verify that T is firmly nonexpansive, has a unique fixed point, and $T(x) = 0$ for all x . Moreover $\Phi = \text{Id} - T = \text{Id}$, which has a constant of metric subregularity $c = 1$. Theorem 1.6 then predicts that the Douglas–Rachford algorithm converges linearly with rate constant 0 in this case, i.e. it converges in one step. The reader can verify that this indeed is the case.

To see the importance of the restriction to the affine subspace W , consider instead of two lines in \mathbb{R}^2 two lines in \mathbb{R}^3 intersecting at the origin. It can be shown that the fixed points of the Douglas–Rachford operator consist of the axis – let’s call it the z axis – extending from the origin, perpendicular to the linear hull of the two lines [6]. It is elementary to verify that, from any starting point x^0 in \mathbb{R}^3 , the Douglas–Rachford algorithm converges in one step to the intersection of the z axis with the affine subspace containing x^0 and parallel to the plane containing the lines A and B . Clearly, the fixed points of the mapping T are not isolated points, but they are isolated points relative to the affine subspace containing x^0 and parallel to A and B , so Theorem 1.6 applies and predicts, correctly, that the Douglas–Rachford algorithm converges to a fixed point in one step. The projection of this fixed point onto the set B is the solution to the problem of finding the point of intersection.

Corollary 1.9 (Polyhedrality implies linear convergence). Let $W \subset V$ be an affine subspace and $T : W \rightrightarrows W$ be quasi-firmly nonexpansive on W . Let $\text{Fix } T \cap W$ be an isolated point, $\{\bar{x}\}$. If T is polyhedral, then there is a neighborhood \mathcal{O} of \bar{x} such that

$$\text{dist}(x_+, \text{Fix } T) \leq \sqrt{1 - \kappa} \text{dist}(x, \text{Fix } T) \quad \forall x_+ \in Tx, \forall x \in \mathcal{O} \cap W,$$

where $0 < \kappa = c^{-2}$ for c a constant of metric subregularity of $\text{Id} - T$ at \bar{x} for the neighborhood $\mathcal{O} \cap W$. Consequently, the fixed point iteration $x^{k+1} = Tx^k$ converges linearly to $\text{Fix } T$ with rate $\sqrt{1 - \kappa}$ for all $x^0 \in \mathcal{O} \cap W$.

Proof. The result follows immediately from Proposition 1.2 and Theorem 1.6. \square

The requirement that the fixed point set is a singleton can be viewed as a uniqueness assumption, which is common in the inverse problems literature. It is well known, however, that, even if the solution to a given problem is unique, the set of fixed points of the numerical method (of interest to us, the Douglas–Rachford operator) need not be solutions to the given problem, much less be unique [6, 44]. Recent work has shown, however, that the set of fixed points need only consist of singletons relative to appropriate affine subspaces where the iterates lie [37, 54]. This feature has been exploited in the analysis of the Douglas–Rachford algorithm applied to problems with polyhedral and quadratic structure [43]. Metric (sub)regularity, on the other hand, is one of the central assumptions of well-posedness of inverse problems [20, 39]. Other useful equivalent characterizations of metric subregularity can be found in [20]. Polyhedrality can be quite easy to verify, as we will see below.

2 Linear Convergence of Douglas–Rachford/ Alternating Directions Method of Multipliers

We consider problems in the following format:

$$\underset{u \in U}{\text{minimize}} J(u) + H(Au). \quad (\mathcal{P}')$$

There are many possibilities for solving such problems. We focus our attention on one of the more prevalent methods, the alternating direction method of multipliers, abbreviated as ADMM (primary sources include [22, 23, 29, 32, 55]). This method is one of many *splitting methods* which are the principle approach to handling the computational burden of large-scale, separable problems [15]. ADMM belongs to a class of *augmented Lagrangian methods* whose original motivation was to regularize Lagrangian formulations of constrained optimization problems.

Introducing a new variable $v \in V$, our problem is to solve

$$\underset{(u,v) \in U \times V}{\text{minimize}} J(u) + H(v), \quad \text{subject to } Au = v. \quad (2.1)$$

The augmented Lagrangian \tilde{L} for (2.1) is given by

$$\tilde{L}(u, v, b) = J(u) + H(v) + \langle b, Au - v \rangle + \frac{\eta}{2} \|Au - v\|^2, \quad (2.2)$$

where $b \in V$, $\eta > 0$ is a fixed penalty parameter. The ADMM algorithm for solving (2.1) is, given (u^k, v^k, b^k) , $k \in \mathbb{N}$, compute $(u^{k+1}, v^{k+1}, b^{k+1})$ by

$$u^{k+1} \in \operatorname{argmin}_u \left\{ J(u) + \frac{\eta}{2} \|Au - v^k + \eta^{-1}b^k\|^2 \right\}; \quad (2.3)$$

$$v^{k+1} \in \operatorname{argmin}_v \left\{ H(v) + \frac{\eta}{2} \|Au^{k+1} - v + \eta^{-1}b^k\|^2 \right\}; \quad (2.4)$$

$$b^{k+1} = b^k + \eta(Au^{k+1} - v^{k+1}). \quad (2.5)$$

Using $\frac{\eta}{2} \|Au - v + \eta^{-1}b^k\|^2 - \frac{1}{2\eta} \|b^k\|^2 = \langle b^k, Au - v \rangle + \frac{\eta}{2} \|Au - v\|^2$, the algorithm (2.3)-(2.5) can be written equivalently as

Algorithm 2.1 (ADMM).

Initialization. Choose $\eta > 0$ and $(v^0, b^0) \in U \times V \times V$.

General Step ($k = 0, 1, \dots$)

$$u^{k+1} \in \operatorname{argmin}_u \left\{ J(u) + \langle b^k, Au \rangle + \frac{\eta}{2} \|Au - v^k\|^2 \right\}; \quad (2.6a)$$

$$v^{k+1} \in \operatorname{argmin}_v \left\{ H(v) - \langle b^k, v \rangle + \frac{\eta}{2} \|Au^{k+1} - v\|^2 \right\}; \quad (2.6b)$$

$$b^{k+1} = b^k + \eta(Au^{k+1} - v^{k+1}). \quad (2.6c)$$

The penalty parameter η need not be a constant, and indeed evidence indicates that the choice of η can greatly impact the complexity of the algorithm, but this is beyond the scope of this investigation, so we have left this parameter fixed.

We do not specify how the argmin in steps (2.6a)-(2.6b) should be calculated, and indeed, the analysis that follows assumes that these can be computed exactly. This

is, of course, not true in practice. In an attempt to circumvent this fact, the standard approach in numerical analysis is to accommodate summable errors. The generalization to summable errors is, however, tantamount to *eventual* exact evaluation of (2.6a)-(2.6b) and thus, for all practical purposes, is no different from *immediate* exact evaluation, the latter involving errors that sum to zero.

Even if we do assume infinite precision, a few remarks about the computational complexity of the individual steps of Algorithm 2.1 are warranted. Inspection of (2.6a) shows that an implicit method involving computation of the inverse of $A^T A$ may not be feasible if this is very large or does not otherwise enjoy a structure that allows for efficient inversion. If J is smooth, a number of classical quasi-Newton methods, with error bounds, are available [49]. If J is nonsmooth, then a forward-backward-type method such as FISTA [10] could be applied. In the latter case new results on convergence of the iterates to a solution open the door to error bounds at this stage [3]. The second step (2.6b) does not involve any matrix inversion, but will, for exact penalization, involve a nonsmooth penalty H . Again, one has recourse to fast first-order methods that, as of very recently, permit error bounds.

Our goal is to determine the rate of convergence of these algorithms so that they may be used as inner routines in an iteratively regularized procedure. Knowing that an algorithm converges linearly, for instance, yields rational stopping criteria with computable estimates for the distance of the current iterate to the solution set.

We present sufficient conditions for *linear* convergence of Algorithm 2.1 by showing the same for the Douglas-Rachford algorithm which is more amenable to the tools of abstract fixed point theory presented in Section 1.2. It is well known [22, 29] that the ADMM algorithm can be derived from the Douglas–Rachford algorithm, and vice versa, and therefore sufficient conditions for convergence of Douglas–Rachford also apply here. The first convergence result for Douglas–Rachford is due to Lions and Mercier [44], under the assumption of strong convexity and Lipschitz continuity of J . Recent published work in this direction includes [30, 31, 34]. Convergence rates with respect to objective values under various assumptions on the objective, all of which involving strong convexity, was established in [34, 50] which is conservative. Local linear convergence of the iterates to a *solution* was established in [14] for linear and quadratic programs using spectral analysis. In the first main result, Theorem 2.3, we describe two conditions that guarantee linear convergence of the ADMM iterates to a solution. The first of these conditions follows from classical results of Lions and Mercier [44]. The second condition is based on work of more recent vintage [36], is much more prevalent in applications and generalizes the results of [14].

The (Fenchel-Legendre) dual problem corresponding to the problem (\mathcal{P}') is (see, for instance [12])

$$\min_{w \in V} J^*(A^T w) + H^*(-w).$$

Here J^* and H^* are the Fenchel conjugates of J and H respectively. Instead of working with this dual, we work with the following equivalent form with the change of variable $v = -w$:

$$\min_{v \in V} J^*(-A^T v) + H^*(v). \quad (\mathcal{D}')$$

Under the assumption that the solutions \bar{u} and \bar{b} of the primal and dual problems exist and that the dual gap is zero, the following two inclusions characterize the solutions of the problems (\mathcal{P}') and (\mathcal{D}') respectively:

$$\begin{aligned} 0 &\in \partial J(\bar{u}) + \partial(H \circ A)(\bar{u}); \\ 0 &\in \partial(J^* \circ (-A^T))(\bar{b}) + \partial H^*(\bar{b}). \end{aligned}$$

In both cases, one has to solve an inclusion of the form

$$0 \in (B + D)(x), \quad (2.7)$$

for general set-valued mappings B and D . For any $\eta > 0$, the Douglas–Rachford algorithm [21, 44] for solving (2.7) is given by

$$b^{k+1} \in T'b^k \quad (k \in \mathbb{N}), \quad (2.8)$$

$$\text{for } T' := \mathcal{J}_{\eta D}(\mathcal{J}_{\eta B}(\text{Id} - \eta D) + \eta D), \quad (2.9)$$

where $\mathcal{J}_{\eta D}$ and $\mathcal{J}_{\eta B}$ are the *resolvents* of ηD and ηB respectively. The connection between the ADMM algorithm (2.6a)-(2.6c) and the Douglas–Rachford algorithm (2.8) was first discovered by Gabay [29] and is derived for convenience in the Appendix.

Given b^0 and $v^0 \in Db^0$, following [57], define the new variable $x^0 := b^0 + \eta v^0$ so that $b^0 = \mathcal{J}_{\eta D}x^0$. We thus arrive at an alternative formulation of the Douglas–Rachford algorithm (2.8):

$$x^{k+1} \in Tx^k \quad (k \in \mathbb{N}), \quad (2.10)$$

$$\text{for } T := \frac{1}{2}(R_{\eta B}R_{\eta D} + \text{Id}) = \mathcal{J}_{\eta B}(2\mathcal{J}_{\eta D} - \text{Id}) + (\text{Id} - \mathcal{J}_{\eta D}), \quad (2.11)$$

where $R_{\eta D}$ and $R_{\eta B}$ are the reflectors of the respective resolvents. This is exactly the form of Douglas–Rachford considered in [44].

Remark 2.1 (proximal mappings of convex functions). *Note that for our application*

$$B := \partial(J^* \circ (-A^T)) \quad \text{and} \quad D := \partial H^*, \quad (2.12)$$

and so the resolvent mappings are the proximal mappings of the convex functions $(J^* \circ (-A^T))$ and H^* respectively, and hence the resolvent mappings and corresponding fixed point operator T are single-valued [47].

Proposition 2.2. *Let $J : U \rightarrow \mathbb{R} \cup \{+\infty\}$ and $H : V \rightarrow \mathbb{R}$ be proper, lsc and convex. Let $A : U \rightarrow V$ be linear and suppose there exists a solution to $0 \in (B + D)(x)$ for B and D defined by (2.12). For fixed $\eta > 0$, given any initial points x^0 and $(b^0, v^0) \in \text{gph } D$ such that $x^0 = b^0 + \eta v^0$, the sequences $(b^k)_{k \in \mathbb{N}}$, $(x^k)_{k \in \mathbb{N}}$ and $(v^k)_{k \in \mathbb{N}}$ defined respectively by (2.8), (2.10) and $v^k := \frac{1}{\eta}(x^k - b^k)$ converge to points $\bar{b} \in \text{Fix } T'$, $\bar{x} \in \text{Fix } T$ and $\bar{v} \in D(\text{Fix } T')$. The point $\bar{b} = \mathcal{J}_{\eta D}\bar{x}$ is a solution to (\mathcal{D}') , and $\bar{v} = \frac{1}{\eta}(\bar{x} - \bar{b}) \in D\bar{b}$. If, in addition, A has full column rank, then the sequence $(b^k, v^k)_{k \in \mathbb{N}}$ corresponds exactly to the sequence of points generated in steps (2.6b) and (2.6c) of Algorithm 2.1 and the sequence $(u^{k+1})_{k \in \mathbb{N}}$ generated by (2.6a) converges to \bar{u} , a solution to (\mathcal{P}') .*

Proof. Following [22, 57], we rewrite the Douglas–Rachford iteration 2.8 in two steps: Given $(b^0, v^0) \in \text{gph } D$, for $k \in \mathbb{N}$ do

$$\text{find } (q^{k+1}, s^{k+1}) \in \text{gph}(B) \text{ such that } q^{k+1} + \eta s^{k+1} = b^k - \eta v^k; \quad (2.13a)$$

$$\text{find } (b^{k+1}, v^{k+1}) \in \text{gph}(D) \text{ such that } b^{k+1} + \eta v^{k+1} = q^{k+1} + \eta v^k. \quad (2.13b)$$

The existence and uniqueness in the above steps follows from the representation lemma [22, Corollary 3.6.3]. The mappings B, D are maximal monotone operators as the subdifferentials of proper lsc convex functions. This together with the fact that the solution set of (2.7) is non-empty yields that the sequence $(b^k, v^k)_{k \in \mathbb{N}}$ defined by the algorithm (2.13) converges to some (\bar{b}, \bar{v}) such that $\bar{v} \in D\bar{b}$ and \bar{b} solves (\mathcal{D}') [57, Theorem 1]. By the change of variables $x^k = b^k + \eta v^k$, it follows that $x^k \rightarrow \bar{x} \in \text{Fix } T$ for T given by (2.11).

For these definitions of B and D , the sequence $(b^k)_{k \in \mathbb{N}}$ generated by $b^k := \mathcal{J}_{\eta D}x^k$ for x^k generated by (2.10) corresponds exactly to the sequence $(b^k)_{k \in \mathbb{N}}$ generated by (2.8). Moreover, if A is full column rank, then by the discussion in [22] (see the Appendix) both $(b^k)_{k \in \mathbb{N}}$ and the sequence $(v^k)_{k \in \mathbb{N}}$ generated by $v^k := \frac{1}{\eta}(x^k - b^k) \in Db^k$ correspond exactly to the sequences of points b^k and v^k generated by (2.6a)-(2.6c). Consequently, by [22, Proposition 3.42]¹ the sequence $(u^k)_{k \in \mathbb{N}}$ defined by (2.6a) converges to a solution of (\mathcal{P}') . \square

We now state sufficient conditions guaranteeing linear convergence of the ADMM and the Douglas–Rachford algorithms. The first conditions (i) of Theorem 2.3 are classical. The second conditions are new.

Theorem 2.3 (local linear convergence I). *Let $J : U \rightarrow \mathbb{R} \cup \{+\infty\}$ and $H : V \rightarrow \mathbb{R}$ be proper, lsc and convex. Suppose there exists a solution to $0 \in (B + D)(x)$ for B and D defined by (2.12) where $A : U \rightarrow V$ is an injective linear mapping. Let $\hat{x} \in \text{Fix } T$ for T defined by (2.11). For fixed $\eta > 0$ and any given triplet of points (b^0, v^0, x^0) satisfying $x^0 := b^0 + \eta v^0$, with $v^0 \in Db^0$, generate the sequence $(v^k, b^k)_{k \in \mathbb{N}}$ by (2.6a)-(2.6c) and the sequence $(x^k)_{k \in \mathbb{N}}$ by (2.10).*

(i) *Let $\mathcal{O} \subset U$ be a neighborhood of \hat{x} on which H is strongly convex with constant μ and ∂H is β -inverse strongly monotone for some $\beta > 0$. Then, for any $(b^0, v^0, x^0) \in \mathcal{O}$ satisfying $x^0 := b^0 + \eta v^0 \in \mathcal{O}$, the sequences $(x^k)_{k \in \mathbb{N}}$ and $(v^k, b^k)_{k \in \mathbb{N}}$ converge linearly to the respective points $\bar{x} \in \text{Fix } T$ and (\bar{b}, \bar{v}) with rate at least $K = (1 - \frac{2\eta\beta\mu^2}{(\mu+\eta)^2})^{\frac{1}{2}} < 1$.*

(ii) *Suppose that $T : W \rightarrow W$ for some affine subspace $W \subset U$ with $\hat{x} \in W$. On the neighborhood \mathcal{O} of \hat{x} relative to W , that is $\mathcal{O} \cap W$, suppose there is a constant $\kappa > 0$ such that*

$$\|x - x^+\| \geq \sqrt{\kappa} \text{dist}(x, \text{Fix } T) \quad \forall x \in \mathcal{O} \cap W, \quad \forall x^+ \in Tx. \quad (2.14)$$

Then the sequences $(x^k)_{k \in \mathbb{N}}$ and $(v^k, b^k)_{k \in \mathbb{N}}$ converge linearly to the respective points $\bar{x} \in \text{Fix } T \cap W$ and (\bar{b}, \bar{v}) with rate bounded above by $\sqrt{1 - \kappa}$.

¹By convergence of $v^k \rightarrow \bar{v}$ and $b^k \rightarrow \bar{b}$ and the update rule (2.6c), $Au^k \rightarrow \bar{v}$, from which the claim follows – see the Appendix.

In either case, the limit point $\bar{b} = \mathcal{J}_{\eta D} \bar{x}$ is a solution to (\mathcal{D}') , $\bar{v} \in D\bar{b}$ and the sequence $(u^k)_{k \in \mathbb{N}}$ given by (2.6a) of Algorithm 2.1 converges to \bar{u} , a solution of (\mathcal{P}') .

Proof. The final statement of the theorem and the statements about the sequence (b^k, v^k) follows from Proposition 2.2 where it is shown that the sequence $(v^k, b^k)_{k \in \mathbb{N}}$ generated by (2.6a)-(2.6c) corresponds to sequences $(b^k)_{k \in \mathbb{N}}$ and $(v^k)_{k \in \mathbb{N}}$ generated respectively by (2.8) and $v^k = \frac{1}{\eta}(x^k - b^k) \in Db^k$ for $(x^k)_{k \in \mathbb{N}}$ generated by (2.10). The linear convergence of the iterates of Algorithm 2.1 claimed in statements (i) and (ii) follows from the properties of the operators T' and T defined respectively by (2.9) and (2.11).

Part (i). Since H is assumed to be strongly convex with $\mu > 0$ the modulus of convexity on \mathcal{O} , ∂H is strongly monotone with modulus of monotonicity μ [5, Example 22.3]. Since ∂H is also maximally monotone, using the identity $\partial H = (\partial H^*)^{-1}$ (see, for example, [53, Corollary 3.49]) we conclude that ∂H^* is Lipschitz continuous with constant $\frac{1}{\mu}$. Moreover, since ∂H is β -inverse strongly monotone on \mathcal{O} , we have for any $x, y \in \mathcal{O}$

$$\langle u - v, x - y \rangle \geq \beta \|u - v\|^2, \quad \text{whenever } u \in \partial H(x), v \in \partial H(y).$$

Hence ∂H^* is strongly monotone with modulus β and Proposition 4 of [44] applies to yield linear convergence of the sequences (x^k) and (b^k) to the respective limit points \bar{x} and \bar{b}

$$\|x^k - \bar{x}\| \leq LK^k; \quad \|b^k - \bar{b}\| \leq LK^k, \quad (2.15)$$

where L is some constant, $K = (1 - \frac{2\eta\beta}{(1+\eta\xi)^2})^{\frac{1}{2}}$ and $\xi = \frac{1}{\mu}$ is the Lipschitz constant for the set-valued map ∂H^* on \mathcal{O} . Now, since $v^k = \frac{1}{\eta}(x^k - b^k)$, we have for $v^k \rightarrow \bar{v} := \frac{1}{\eta}(\bar{x} - \bar{b})$ with the same rate as x^k and b^k , modulo a constant:

$$\|v^k - \bar{v}\| \leq \frac{1}{\eta} (\|x^k - \bar{x}\| + \|\bar{b} - b^k\|) \leq \frac{2LK^k}{\eta}. \quad (2.16)$$

This completes the proof of the first statement. \triangle

Part (ii). Since B and D are maximal monotone operators the reflected resolvents $R_{\eta B}$ and $R_{\eta D}$ are nonexpansive [5, Proposition 23.7]. The composition $R_{\eta B}R_{\eta D}$ is nonexpansive which implies that the mapping T is firmly nonexpansive [5, Proposition 4.2], and hence quasi-firmly nonexpansive on W . Condition (2.14) is the coercivity condition (b) of [36, Lemma 3.1] which guarantees local linear convergence of fixed-point iterations for (S, ϵ) -firmly nonexpansive mappings ($S \subset \text{Fix } T \cap W$). Quasi-firmly nonexpansive mappings, under consideration here, are $(\text{Fix } T \cap W, 0)$ -firmly nonexpansive. Thus, by [36, Lemma 3.1] the sequence $(x^k)_{k \in \mathbb{N}}$ converges linearly on the neighborhood \mathcal{O} with rate $\sqrt{1 - \kappa}$. Nonexpansiveness of the resolvent $\mathcal{J}_{\eta D}$ and the relations $b^k = \mathcal{J}_{\eta D}x^k$ and $v^k = \frac{1}{\eta}(x^k - b^k)$ then complete the proof of the second statement. \square

Remark 2.4. The strong convexity assumption (i) of Theorem 2.3 fails in a wide range of applications, and in particular for feasibility problems (minimizing the sum

of indicator functions). By Theorem 1.6, case (ii) of Theorem 2.3, in contrast, holds in general for mappings T for which $\text{Id} - T$ is metrically subregular and the fixed point sets are isolated points with respect to an affine subspace to which the iterates are confined. The restriction to the affine subspace W is a natural generalization for the Douglas–Rachford algorithm, where the iterates are known to stay confined to affine subspaces orthogonal to the fixed point set [37, 54]. It would be far too restrictive to require that $\text{Fix } T$ be a singleton on the entire ambient space V rather than with respect to just the affine hull of the iterates. We show that metric subregularity with respect to this affine subspace holds in many applications. (See also Example 1.8.)

Remark 2.5. Proposition 2.2 and Theorem 2.3 and their proofs also hold in infinite dimensional Hilbert spaces. Lemma 3.1 of [36] is stated for Euclidean spaces, but the proof holds also on general Hilbert spaces.

Proposition 2.6 (polyhedrality of the Douglas–Rachford operator). *Let $J : U \rightarrow \mathbb{R} \cup \{+\infty\}$ and $H : V \rightarrow \mathbb{R}$ be proper, lsc and convex. Suppose, in addition, that J and H are piecewise linear-quadratic. The operator $T : V \rightarrow V$ defined by (2.11) with $\eta > 0$ fixed, is polyhedral for B and D given by (2.12) where $A : U \rightarrow V$ is a linear mapping.*

Proof. Since the functions J and H are proper, lsc, convex and piecewise linear-quadratic, by [56, Theorem 11.14] so are the Fenchel conjugates, J^* and H^* . The subdifferentials $B := \partial(J^* \circ (-A^T))$ and $D := \partial H^*$ and their resolvents, therefore, are polyhedral mappings [56, Proposition 12.30]. Since the graphs of reflectors $R_{\eta B}$ and $R_{\eta D}$ correspond to the graphs of their respective resolvents $\mathcal{J}_{\eta B}$ and $\mathcal{J}_{\eta D}$ through a linear transformation, $R_{\eta B}$ and $R_{\eta D}$ are also polyhedral mappings. Since by Remark 2.1 the resolvents $\mathcal{J}_{\eta B}$ and $\mathcal{J}_{\eta D}$ are single-valued, the reflectors $R_{\eta B}$ and $R_{\eta D}$ are also single-valued. Therefore $T = \frac{1}{2}(R_{\eta B}R_{\eta D} + I)$ is polyhedral as the composition of single-valued polyhedral mappings. \square

Theorem 2.7 (local linear convergence II). *Let $J : U \rightarrow \mathbb{R} \cup \{+\infty\}$ and $H : V \rightarrow \mathbb{R}$ be proper, lsc, convex, piecewise linear-quadratic functions (see Definition 1.3). Define the operator $T : V \rightarrow V$ by (2.11) with $\eta > 0$ fixed and B and D given by (2.12) where $A : U \rightarrow V$ is a linear mapping. Suppose that there exists a solution to $0 \in (B + D)(x)$, that $T : W \rightarrow W$ for W some affine subspace of V and that $\text{Fix } T \cap W$ is an isolated point $\{\bar{x}\}$. Then there is a neighborhood \mathcal{O} of \bar{x} such that, for all starting points (x^0, v^0, b^0) with $x^0 := b^0 + \eta v^0 \in \mathcal{O} \cap W$ for $v^0 \in D(b^0)$ so that $\mathcal{J}_{\eta D}x^0 = b^0$, the sequence $(x^k)_{k \in \mathbb{N}}$ generated by (2.10) converges linearly to \bar{x} where $\bar{b} := \mathcal{J}_{\eta D}\bar{x}$ is a solution to (\mathcal{D}') . The rate of linear convergence is bounded above by $\sqrt{1 - \kappa}$, where $\kappa = c^{-2} > 0$, for c a constant of metric subregularity of $\text{Id} - T$ at \bar{x} for the neighborhood \mathcal{O} . Moreover, the sequence $(b^k, v^k)_{k \in \mathbb{N}}$ generated by Algorithm 2.1 converges linearly to (\bar{b}, \bar{v}) with $\bar{v} = \frac{1}{\eta}(\bar{x} - \bar{b})$, and the sequence $(u^k)_{k \in \mathbb{N}}$ defined by (2.6a) of Algorithm 2.1 converges to a solution to (\mathcal{P}') .*

Proof. By Proposition 2.6 the Douglas–Rachford operator T is polyhedral and thus the first statement follows from Corollary 1.9. The statement about the sequences generated by Algorithm 2.1 follows as in Theorem 2.3. \square

3 Error Bounds and Iterative Penalization

In this section, we study an iteratively regularized algorithmic scheme for solving the problems of the form

$$\min \{ J(u) \mid u \in U \text{ and } f_j(Au) \leq \epsilon_j, j = 1, 2, \dots, M \},$$

where $J : U \rightarrow (-\infty, +\infty]$ is proper lsc and convex, the mapping $A : U \rightarrow V$ is linear, for all j the nonnegative-valued function $f_j : V \rightarrow \mathbb{R}_+$ is convex and smooth (at least at points that matter) and $\epsilon_j > 0$. We refer to the inequality constraints as *structured constraints*. It will be convenient to introduce the following notation that will help to reduce clutter. We collect the constraints into a vector-valued function so that we can write the problem as

$$\begin{array}{ll} \underset{u \in U}{\text{minimize}} & J(u) \\ \text{subject to} & F_\epsilon(Au) \leq 0, \end{array} \quad (\mathcal{P})$$

where

$$F_\epsilon : V \rightarrow \mathbb{R}^M := v \mapsto (f_1(v) - \epsilon_1, f_2(v) - \epsilon_2, \dots, f_M(v) - \epsilon_M)^T. \quad (3.1)$$

Here the vector inequality is understood as holding element-wise.

A common approach to solving problems of the type (\mathcal{P}) arising from inverse problems is to apply *implicitly* the structured constraint by adding some (usually smooth) quantification of the constraint violation into the objective function:

$$\underset{u \in U}{\text{minimize}} \quad J(u) + \rho \theta(F_\epsilon(Au)), \quad (\mathcal{P}_\rho)$$

where $\theta : \mathbb{R}^M \rightarrow (-\infty, +\infty]$ is a proper, lsc convex function and $\rho > 0$. This places us in the context of the previous section since problem (\mathcal{P}_ρ) is the specialization of (\mathcal{P}') with $H(Au) = \rho \theta(F_\epsilon(Au))$.

As is often seen in the inverse problems literature, the constraint violation parameter $\epsilon_j = 0$ ($j = 1, \dots, M$), essentially penalizing divergence from the origin. A prominent instance of this form of regularization is the squared norm: $\theta(v) := \|v\|^2$. There are many efficient methods available for solving (\mathcal{P}_ρ) . It is clear that for a certain value of ρ the optimal solution to (\mathcal{P}_ρ) , u_ρ , will satisfy $f_j(Au_\rho) \leq \bar{\epsilon}_j(\rho)$ with the *effective* error $\bar{\epsilon}_j(\rho)$ depending on ρ . What is *not* true in general, however, is that the solution to (\mathcal{P}_ρ) corresponds to the solution to (\mathcal{P}) for the constraint error $\bar{\epsilon}(\rho)$. Moreover, for our intended applications, U is a finite dimensional Euclidean space with dimension n and the dimensionality of the constraints M grows superlinearly as a function of n , so we would like to consolidate the constraints somehow while exploiting the phenomenon that, at the solution to (\mathcal{P}) relatively few of the constraints are in fact tight or *active*.

We consider convex penalties that reduce the dimensionality of the constraint structure and have the property that $\theta(F_\epsilon(Au)) = 0$ if and only if $F_\epsilon(Au) \leq 0$. Of particular interest among penalties with this property are *exact penalties*, that is penalties θ with the property that solutions to (\mathcal{P}_ρ) correspond to solutions to (\mathcal{P}) for

all values of ρ beyond a certain threshold $\bar{\rho}$. For more background on exact penalization see, for example, [11, 17, 19, 25, 33, 46]. We point also to Friedlander and Tseng [28] for a connection between exact penalization and what they call *exact regularization* as this fits well with our viewpoint that the structured constraints $F_\epsilon(Au) \leq 0$ constitute a regularization of the *model* with regularization parameter ϵ . This illustrates the distinction between *model-based* regularization, that is, regularization of the constraints motivated by external (eg. statistical) considerations, versus *numerical* regularization motivated solely on the grounds of enabling efficient (approximate) numerical solutions to (\mathcal{P}) .

While it is nice to know that, with exact penalization, one can achieve an exact correspondence between the original constrained optimization problem and the penalized problem, the whole point of relaxing the constraints is to reduce the computational burden of strictly enforcing the constraints. As is often done in practice, one gradually strengthens the constraints, finding intermediate points that nearly solve the relaxed problem and using these as starting points for solving a more strictly penalized problem. Together with Theorem 3.4 below, the linear convergence rate established in Theorems 2.3 and 2.7 of the previous section yield estimates on the distance of intermediate points to the solution set of the relaxed problem as well as estimates on the distance to feasibility for the unrelaxed problem.

3.1 Structured Constraints and penalization

Define

$$\mathcal{C} := \{u \in U \mid F_\epsilon(Au) \leq 0\}. \quad (3.2)$$

This is a closed convex set since the f_j are lsc and convex. If there exists some $\alpha \in \mathbb{R}$ such that $\mathcal{C} \cap \text{lev}_{\leq \alpha} J$ is nonempty and bounded then (\mathcal{P}) has a solution [5, Theorem 11.9]. This will happen, for instance, if $\text{dom}(J) \cap \mathcal{C} \neq \emptyset$ and J is *coercive* [5, Proposition 11.12], that is J satisfies

$$\lim_{\|u\| \rightarrow \infty} J(u) = +\infty. \quad (3.3)$$

Such assumptions are naturally satisfied in many applications. Moreover, $\text{lev}_{\leq \bar{\alpha}} J(u)$, the lower level-set of J corresponding to the optimal value $\bar{\alpha}$ in (\mathcal{P}) , is convex and so the set of optimal solutions to (\mathcal{P}) is also convex. Define $J_\rho := J + \rho \theta(F_\epsilon \circ A)$ for the convex, lsc function θ satisfying $\theta(w) \geq 0$ for all w and $\theta(w) = 0$ if and only if $F_\epsilon(w) \leq 0$. Then J_ρ is convex, lsc and corresponds exactly to J on the set \mathcal{C} . Otherwise J_ρ increases pointwise to $+\infty$ at points outside \mathcal{C} as $\rho \rightarrow \infty$. For $(\rho_k)_{k \in \mathbb{N}}$ with $\rho_k \rightarrow \infty$, the sequence of functions (J_{ρ_k}) epi-converges (see [56, Definition 7.1]) to $J + \iota_{\mathcal{C}}$ as $k \rightarrow +\infty$ where $\iota_{\mathcal{C}}$ is the indicator function of the set \mathcal{C} . As we will allow *approximate* solution of problems (\mathcal{P}_ρ) it will be helpful to recall the set of γ -minimizers: $\gamma - \text{argmin} J_\rho := \{u \mid J_\rho(u) \leq \inf J_\rho + \gamma\}$. The relation between the solution sets to (\mathcal{P}) and (\mathcal{P}_ρ) is detailed in the following, which is a direct application of [56, Theorem 7.33].

Proposition 3.1. Let $J : U \rightarrow (-\infty, +\infty]$, $F_\epsilon : V \rightarrow \mathbb{R}^M$ and $\theta : \mathbb{R}^M \rightarrow \mathbb{R}$ be proper, lsc and convex, and let $A : U \rightarrow V$ be linear. Let J be coercive with $\text{dom } J \cap \mathcal{C} \neq \emptyset$ for \mathcal{C} defined by (3.2). Suppose further that $\theta(w) \geq 0$ and that $\theta(w) = 0$ if and only if $F_\epsilon(w) \leq 0$. Define $J_{\rho_k} := J + \rho_k \theta(F_\epsilon \circ A)$ where $\rho_k \nearrow +\infty$ as $k \nearrow +\infty$. Then $\inf J_{\rho_k} \rightarrow \inf J + \iota_{\mathcal{C}} < +\infty$. Moreover, for any sequence of errors $\gamma_k \searrow 0$ and corresponding points $u^k \in \gamma_k - \text{argmin } J_{\rho_k}$, the sequence $(u^k)_{k \in \mathbb{N}}$ is bounded, and all its cluster points belong to $\text{argmin } \{J + \iota_{\mathcal{C}}\}$.

Proof sketch. The property of the convex penalty θ that $\theta(w) \geq 0$ and $\theta(w) = 0$ if and only if $F_\epsilon(w) \leq 0$ yields epi-convergence of J_{ρ_k} to $J + \iota_{\mathcal{C}}$. Coercivity of J guarantees that J_ρ is level bounded for all values of $\rho > 0$. These two properties, together with lower semicontinuity and the fact that J and J_ρ are proper, are all that is needed to prove the result. \square

If the regularization were *exact*, then we would know that for all parameter values ρ large enough, the solutions to (\mathcal{P}_ρ) coincide with solutions to (\mathcal{P}) . We return to this later.

3.2 Solution to the regularized Subproblem and error bounds

We now turn our attention to solution of the problem (\mathcal{P}_ρ) for a fixed value of ρ_k . The ADMM algorithm discussed in Section 2 is useful for solving this problem in the sense that it has an error bound under specific assumptions which gives a stopping rule. This is not unique to Algorithm 2.1, but we focus on this method due to its prevalence in practice.

Recall the exact problem (\mathcal{P}) :

$$\begin{aligned} & \underset{u \in U}{\text{minimize}} && J(u) \\ & \text{subject to} && F_\epsilon(Au) \leq 0. \end{aligned} \tag{\mathcal{P}}$$

It will be convenient to rewrite the penalized problem² (\mathcal{P}_ρ) as

$$\underset{u \in U}{\text{minimize}} \frac{1}{\rho} J(u) + \theta(F_\epsilon(Au)). \tag{\mathcal{P}_\rho}$$

Consider also the *limiting* problem

$$\underset{u \in U}{\text{minimize}} \theta(F_\epsilon(Au)). \tag{\mathcal{P}_\infty}$$

We view problem (\mathcal{P}_ρ) as the regularized version of (\mathcal{P}_∞) with J as the regularizing functional and $\frac{1}{\rho}$ as the regularization parameter. Denote the solution sets to these problems by

$$\begin{aligned} S &:= \text{argmin } \{J(u) \mid u \in U, F_\epsilon(Au) \leq 0\}, \\ S_\rho &:= \text{argmin} \left\{ \frac{1}{\rho} J(u) + \theta(F_\epsilon(Au)) \mid u \in U \right\}, \\ S_\infty &:= \text{argmin } \{\theta(F_\epsilon(Au)) \mid u \in U\}. \end{aligned}$$

²Of course, the value of the problem is not the same, but the solutions are.

If the penalization θ satisfies $\theta(F_\epsilon(Au)) = 0$ if and only if $F_\epsilon(Au) \leq 0$, then it is immediately clear that S_∞ corresponds to the feasible set of problem (\mathcal{P}) hence $S \subset S_\infty$. What is more remarkable is that, if a Lagrange multiplier for (\mathcal{P}) exists, then $S_\rho = S$ for all ρ large enough, that is, the penalty θ is *exact*.

Theorem 3.2 (Theorem 4.2 of [28]). *Suppose that S is nonempty and compact, and that there exist Lagrange multipliers λ for (\mathcal{P}) . Let the penalization θ in (\mathcal{P}_ρ) be convex. Assume, moreover, that θ satisfies the condition $\theta(F_\epsilon(Au)) = 0$ if and only if $F_\epsilon(Au) \leq 0$. Then the solution set to the penalized problem, S_ρ , coincides with the solution set to the exact problem, S , for all $\rho > \theta^\circ(\lambda)$ where θ° is the polar function of θ given by $\theta^\circ(\lambda) = \sup_{x \notin S} \frac{\lambda^T x}{\theta(x)}$.*

It is easy to check whether a solution $u_\rho \in S_\rho$ is in fact feasible for (\mathcal{P}) (and hence also in S) by simply evaluating the value of $\theta(F_q(Au_\rho))$. More generally, one would check whether the first order optimality conditions for (\mathcal{P}_∞) are satisfied at u_ρ , namely

$$0 \stackrel{?}{\in} \partial\theta(F_\epsilon \circ A(\cdot)) \text{ at } u_\rho. \quad (3.4)$$

An explicit formula for the subdifferential in (3.4) for image denoising and deconvolution is given in Section 4 as this will be needed for computing Step (2.6b) of Algorithm 2.1.

If, in addition, S_∞ is *weakly sharp* (see Definition 1.4), then one can obtain an upper bound for the distance of solutions to (\mathcal{P}_ρ) to *feasible* solutions to (\mathcal{P}) , even in the absence of Lagrange multipliers for (\mathcal{P}) .

Assumption 3.3.

- (i) The solution set S_∞ of problem (\mathcal{P}_∞) is nonempty.
- (ii) $\text{lev}_{\leq \alpha} J$ is bounded for each $\alpha \in \mathbb{R}$ and $\inf_{x \in U} > -\infty$.
- (iii) The solution set S_∞ of (\mathcal{P}_∞) is weakly sharp of order $\nu \geq 1$.

Theorem 3.4. *Suppose Assumption 3.3(i)-(ii) hold.*

- (i) For any $\bar{\rho} > 0$, $\bigcup_{\rho \geq \bar{\rho}} S_\rho$ is bounded.
- (ii) If, in addition, Assumption 3.3(iii) holds with modulus of sharpness ν , then for any $\bar{\rho} > 0$ there exists $\tau > 0$ such that

$$\text{dist}(u_\rho, S_\infty)^{\nu-1} \leq \frac{\tau}{\rho}, \quad \forall u_\rho \in S_\rho, \quad \rho \geq \bar{\rho}. \quad (3.5)$$

- (iii) If, in addition, Assumption 3.3(iii) holds and the penalization θ is exact, then for all ρ large enough, $u_\rho \in S$ and $\text{dist}(u_\rho, S) = \text{dist}(u_\rho, S_\infty) = 0$.

Proof. (i) and (ii). Under the assumption 3.3, Theorem 5.1 in [28] directly applies to yield the result. \triangle

(iii). If the penalization θ is exact, then $\theta(F_\epsilon(Au)) = 0$ if and only if $F_\epsilon(Au) \leq 0$, hence $S = S_\rho$ for all ρ large enough, and S_∞ corresponds exactly to the feasible set in (\mathcal{P}) . \square

Remark 3.5. The error bound (3.5) holds independent of the existence of Lagrange multipliers for (\mathcal{P}) , hence, for exact penalization under Assumption 3.3, Theorem 3.4 yields an upper bound on the distance of solutions to (\mathcal{P}_ρ) to feasible points for (\mathcal{P}) .

4 Application: image deconvolution and denoising with statistical multiscale analysis

We specialize the above results to the application of optimization with statistical multiscale side constraints. All of the examples considered in this section satisfy the requirements of Theorem 2.7, and thus for each fixed value of the penalty parameter ρ local linear convergence to a solution of (\mathcal{P}_ρ) is guaranteed. Moreover, the penalty function θ that we use is *exact* and hence by Theorem 3.2, for ρ large enough, the computed solution to (\mathcal{P}_ρ) is also a solution to (\mathcal{P}) . What is not known *a priori* is what value of ρ yields the correspondence. Moreover, since the whole point of the relaxation (\mathcal{P}_ρ) is to remove the burden of satisfying the constraints, we approach a solution to (\mathcal{P}) via a sequence of solutions to (\mathcal{P}_ρ) for progressively larger values of ρ . This is described precisely in the following sequentially penalized algorithm.

Algorithm 4.1 (Exactly Penalized Sequential ADMM).

Initialization. Given an image y , a sequence of error tolerances $(\gamma_k)_{k \in \mathbb{N}}$ with $0 \leq \gamma_k \rightarrow 0$. Choose parameters: $\beta > 1$ and the penalty parameter $\eta \in (0, 2)$. Initialize $k = i = 0$, $b^{(0,0)} = 0$, $v^0 = y$, $u^{(0,0)} = A^T y$, and compute $u^{(0,1)} = \operatorname{argmin}_u \{ J(u) + \langle b^{(0,0)}, Au \rangle + \frac{\eta}{2} \|Au - v^{(0,0)}\|^2 + \frac{1}{2} \|u - u^{(0,0)}\|^2 \}$.

For $k = 0, 1, 2, \dots$

- **While** $\|u^{(k,i+1)} - u^{(k,i)}\| > \gamma_k$
 - Compute $(v^{(k,i+1)}, b^{(k,i+1)})$ via Algorithm 2.1 steps (2.6b)-(2.6c) with $H := \rho_k \theta(F_q(\cdot))$ for the exact penalty θ and structured constraints F_q .
 - Increment $i = i + 1$ and calculate $u^{(k,i+1)}$ via Algorithm 2.1 step (2.6a).
- **Update/reset:** Set $u^{(k+1,1)} := u^{(k,i+1)}$ and $\rho_{k+1} = \beta \rho_k$. Set $k = k + 1$ and $i = 0$. If $\theta(F_q(u^{(k,1)})) = 0$, set $\gamma_k = 0$.

The outer iteration, indexed by k , consists of numerical approximations to solutions of (\mathcal{P}_ρ) for the penalty parameter ρ_k . The inner iteration proceeds with the current value of ρ_k until the step size between successive iterates $u^{k,i+1}$ and $u^{k,i}$ drops in a linear fashion below a given tolerance γ_k . From Theorem 2.7 one can then obtain a posteriori estimates on the distance of the iterate $u^{k+1,j}$ to the true solution. Then ρ_k is increased by a constant factor. Since, for this model the penalization θ is exact, once the constraints appear to be satisfied (as determined by monitoring the value of $\theta(F_q(v^k))$), it is reasonable to conclude that the correspondence between problems

(\mathcal{P}_ρ) and (\mathcal{P}) holds, and the penalty ρ_k no longer needs to be updated; the inner loop of the algorithm then can be run to the desired accuracy. As indicated in Figures 1(b) and 4, the constraints appear to be satisfied when the penalty term $\rho_k \theta(F_q(v^k))$ (green plot) drops suddenly to machine precision.

The application problem involves image deconvolution and denoising with statistical multiscale estimation as presented in [2, 26, 27]. We are well aware that there are many ways to model such problems that permit much less computationally intensive numerical solutions than the technique we present here. Our interest in multiresolution deconvolution/denoising is two-fold: first, it is one of the few techniques available that has the potential to yield quantitative (i.e. statistical) guarantees for the recovered images, and secondly, it is an important instance of convex optimization problems where the number of constraints grows superlinearly as a function of the number of unknowns. Our numerical demonstration addresses the first issue of quantitative image denoising: if the numerics do not permit estimates for the distance to the model solution, then the quantitative assurances of the model are irrelevant. Unlike the numerical approach proposed in [26, 27], the numerical approach we present here permits error bounds to within machine accuracy of our numerical solution to the true model solution.

Following the approach proposed in [26] we quantify the difference between an estimate $v = Au$ and the data y via the maximum absolute value of all weighted inner products of the residual function $\Delta(\cdot; y) : \mathbb{R}^n \rightarrow \mathbb{R}^n$:

$$f_j(v) := |\langle \omega_j, \Delta(v; y) \rangle|, \quad j \in \{1, 2, \dots, M\}. \quad (4.1)$$

The residual function used in [26] Δ is simply $v - y$. The weights ω_j are scaled window functions so that the set $\mathcal{I} \subset \{1, 2, \dots, M\}$ is the index set corresponding to all collections of these subsets of the image. The statistical multiscale analysis requires that, on each window,

$$\max_{j \in \mathcal{I}} \{f_j(v)\} \leq q. \quad (4.2)$$

The same error q is specified at all scales. Hence F_ϵ in (3.1) specializes to

$$F_q : \mathbb{R}^n \rightarrow \mathbb{R}^{M+1} := v \mapsto (f_1(v) - q, f_2(v) - q, \dots, f_M(v) - q, 0)^T, \quad (4.3)$$

for $f_j : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by (4.1) ($j = 1, \dots, M$) and

$$\theta : \mathbb{R}^{M+1} \rightarrow \mathbb{R} : \theta(w) := \max\{w_1, w_2, \dots, w_{M+1}\}. \quad (4.4)$$

(Here we are expanding the original F_ϵ by the constant function $f_{M+1}(v) := 0$.) The max function is a standard tool in exact penalization methods [17, 19] and falls naturally into the context of piecewise linear-quadratic functions.

Algorithm 4.1 does not specify how the iterates $u^{(k,j)}$ and $v^{(k,j)}$ are calculated. The linear convergence of the inner iterations predicted in Theorem 2.7, from which error bounds can be determined, as well as the numerical convergence of the outer iterates to problem (\mathcal{P}) is discussed next.

4.1 Prox-evaluation

Computation of $u^{(k,i+1)}$ and $v^{(k,i+1)}$ in Algorithm 4.1 involves minimizing the sum of a convex quadratic function and (in general) a convex, nonsmooth, piecewise linear-quadratic function. This can be solved via any number of techniques ranging from first order methods like FISTA [10] to higher-order nonlinear optimization methods like quasi-Newton methods studied in [41]. In order to take advantage of the relative sparsity of the active constraints, we propose the following (exact) algorithm.

Algorithm 4.2 (Steepest Subdifferential Descent).

Initialization. Given b , u , the constant $\eta > 0$ and an initial point v^0 , compute the residual $r^0 := b + \eta Au - \eta v^0$ and the projected residual $z^0 := P_{\partial(\theta(F_q(v^0)))}(r^0)$ for $\partial(\theta(F_q(v^0)))$ given by (4.10).

For $l = 0, 1, 2, \dots$

- **If** $z^l = r^l$
 - set $\bar{v} = v^l$ and *STOP*;
- **else**
 - set $v^{l+1} = v^l + \lambda_l(z^l - r^l)$ where $\lambda_l > 0$ is the largest constant λ such that $\theta(F_q(v^l + \lambda(z^l - r^l))) = f_i(v^l + \lambda(z^l - r^l)) - q$ for $i \in I(v^l)$ with
$$I(v) := \{j \mid f_j(v) - q = \theta(F_q(v))\}; \quad (4.5)$$
 - compute $r^{l+1} := b + \eta Au - \eta v^{l+1}$ and the projected residual
$$z^{l+1} := P_{\partial(\theta(F_q(v^{l+1})))}(r^{l+1}); \quad (4.6)$$
 - increment $l = l + 1$.

Algorithm 4.2 is an *active set method* and the set $I(v)$ defined by (4.5) is the set of *active indexes at v* . Another helpful interpretation is as a steepest subgradient descent method for solving

$$\operatorname{argmin}_v \{G(v) := \theta(F_q(v)) - \langle b, v \rangle + \frac{\eta}{2} \|Au - v\|^2\}. \quad (4.7)$$

The steepest descent step is

$$v^{l+1} = v^l + \lambda_l d^l,$$

for $d^l := P_{\partial G(v)}(0) = -r^l + z^l$ with $z^l := P_{\partial\theta(F_q(v^l))}(r^l)$ and $r^l = b + \eta(Au - v^l)$. The choice of the step length λ_l ensures that, at each step l , the active set is growing; specifically,

$$I(v^l) \subset I(v^l + \lambda_l d^l).$$

At termination, the subdifferential $\partial\theta(F_q(v^l))$ is large enough that it contains the residual r^l . The terminal point of Algorithm 4.2, \bar{v} , is a point in (4.7) since it satisfies the first-order optimality conditions:

$$0 = \bar{z} - b - \eta(Au - \bar{v}) \in \partial\theta(F_q(\bar{v})) - b - \eta(Au - \bar{v}) = \partial G(\bar{v}), \quad (4.8)$$

where $\bar{z} = P_{(\partial\theta(F_q(\bar{v})))}(b + \eta(Au - \bar{v}))$. Replacing u and b with $u^{(k,i+1)}$ and $b^{(k,i)}$ respectively yields the update for $v^{(k,i)}$ in Algorithm 4.1.

The expression for the subdifferential $\partial\theta(F_q)$ is particularly simple in this case. Note that $I(v) \neq \emptyset$ for all v . Applying the (convex) calculus of subdifferentials to the objective $\theta(F_\epsilon(v))$, as permitted by the regularity of θ and F (see, for instance [18, Section 2.3]), yields

$$\partial\theta(F_q(v)) = \text{co } \{\nabla f_j(v) \mid j \in I(v)\}, \quad (4.9)$$

where co denotes the convex hull of a set of points. This, of course, assumes that f_j is differentiable at v for those $j \in I(v)$. Inspection of (4.1) shows that this is not the case in general, in particular at points v^* where $f_j(v^*) = 0$. However, such points will never be in the active set $I(v^*)$ since $f(v^*) - q < 0 \leq \theta(F_q(v^*))$ for all $q > 0$, so we can safely apply formula (4.9) without further ado. This yields the following specialization for $f_j(v) = |\langle w_j, v - y \rangle|$ given by (4.1):

$$\begin{aligned} \partial\theta(F_q(v)) &= \text{co } \{\nabla f_j(v) \mid j \in I(v)\} \\ &= \begin{cases} \text{co } \{\{\text{sign}(\langle w_j, v - y \rangle) w_j \mid j \in I(v) \setminus \{M+1\}\}, 0\} & \theta(F_q(v)) \leq 0 \\ \text{co } \{\text{sign}(\langle w_j, v - y \rangle) w_j \mid j \in I(v)\} & \theta(F_q(v)) > 0. \end{cases} \end{aligned} \quad (4.10)$$

4.2 Synthetic data

Fig. 1 shows a set of synthetic exact data $u^* \in \mathbb{R}^n$ (shown in blue) and corresponding noisy data $y \in \mathbb{R}^n$ (shown in green) with $n = 512$ data points, as well as the reconstructed/denoised signal $\bar{u} \in \mathbb{R}^n$ (shown in red). In this example we consider only denoising, that is, the imaging operator A is the identity so $v = u$. The noisy data y was generated by adding i.i.d. Gaussian random noise with standard deviation $\sigma = 0.05$ to each original data point of u^* . In our specialization of problem (P) we use the total variation penalty

$$J(u) := a \|\nabla u\|_2^2, \quad (4.11)$$

where ∇ is the (discrete) gradient operator. The structured constraints are given by (4.2). The weights $w_j \in \mathbb{R}^n$ are scaled window functions of all intervals of lengths between 1 and 20 pixels, and $\mathcal{I} \subset \{1, 2, \dots, M\}$ is the index set corresponding to all collections of successive pixels in $\{1, 2, \dots, n\}$ of cardinality – or length – from 1 to 20. The same error q is specified at all scales.

For a signal length $n = 512$ with interval lengths from 1 to 20 the number of windows is $M = 10050$. The constant α is, strictly speaking, redundant but was introduced as an additional means to balance the contributions of the individual terms to make the most of limited numerical accuracy (double precision). We chose $\alpha = 0.01$. The constant q was taken to be 2σ .

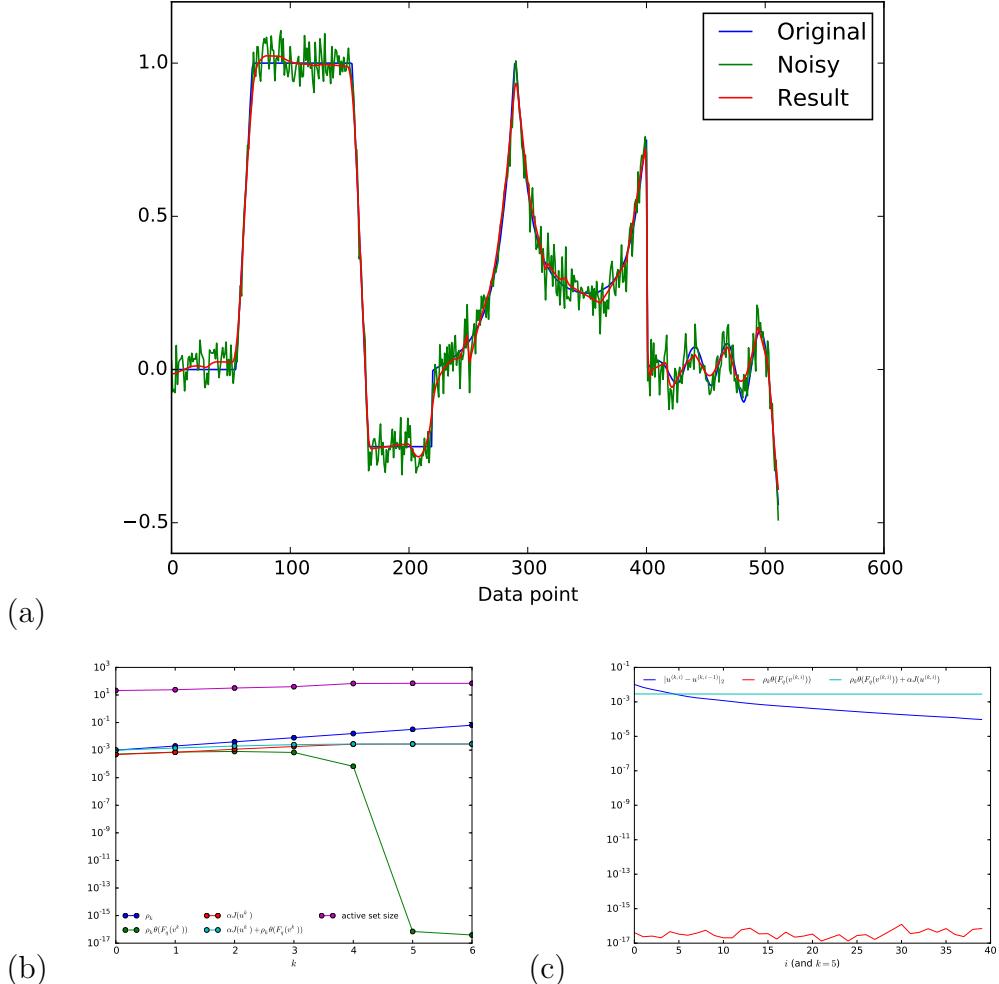


Figure 1: (a) Original, noisy and reconstructed data for a one-dimensional denoising problem. (b) Outer iterates k of Algorithm 4.1 showing solutions, the constraint violation, the active set size and the objective value for the penalized problem (\mathcal{P}_ρ) for successively larger values of the penalty parameter ρ . (c) Inner iterates of Algorithm 4.1 with $\rho_5 = .032$: step sizes, constraint violation, objective value and gap between the primary, domain-space variables $u^{(k,i)}$.

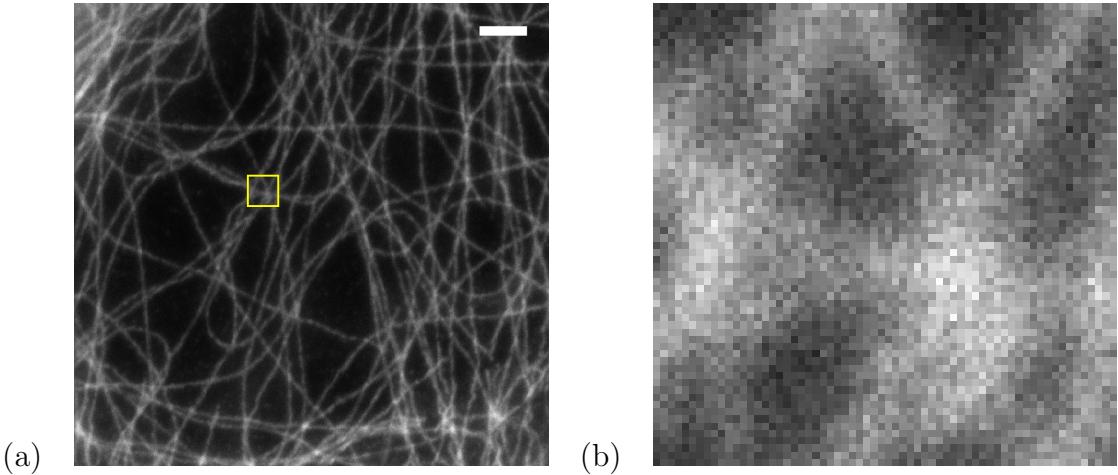


Figure 2: (a) Original data (STED image of Tubulin), (b) an enlargement of the indicated box to be processed. The length of scale bar in (a) is $1\mu\text{m}$, the size of the reconstruction window (b) is $640 \times 640 \text{ nm}^2$.

Figure 1(a) shows very good correspondence of the reconstructed signal to the original. The multi-resolution constraint prevents the usual ‘‘blocky’’ artifacts common to image denoising with TV-regularization. The eventual (starting from around iteration 15) linear convergence of the algorithm can be seen in Figure 1(c). Under the assumption that the latter iterates are indeed in the region of local linear convergence, the observed convergence rate is $c = 0.9245$, which yields an a posteriori upper bound on the distance of the 39th iterate to the true solution: $\|u^{39} - u^*\| \leq \frac{c}{1-c} \|u^{38} - u^{39}\| = 0.001244$. Since the signal length is 512, this amounts to 5 digits of accuracy in the pointwise value of the signal.

4.3 Laboratory data

For our main demonstration, we are presented with an image $y \in \mathbb{R}^n$ (Figure 2(a)) generated from a Stimulated Emission Depletion (STED) microscopy experiment [35, 38] conducted at the Laser-Laboratorium Göttingen examining tubulin, represented as the ‘‘object’’ $u \in \mathbb{R}^m$. The imaging model is simple linear convolution, $Au \approx y$ where A is a convolution matrix with a nonsymmetric experimentally measured point-spread function (290nm^2). The measurement y is *noisy* or otherwise inexact, and thus an exact solution $Au = y$ is not desirable. Although the noise in such images is usually modeled by Poisson noise, a Gaussian noise model with constant variance suffices as the photon counts are of the order of 100 per pixel and do not vary significantly across the image. Figure 2(b) shows a close-up which we used as the noisy data $y \in \mathbb{R}^2$ with $n = 64 \times 64$ data points. We calculate the numerically reconstructed tubulin density \bar{u} shown in Figure 3(a) via Algorithm 4.1 for the problem (\mathcal{P}_ρ) with the qualitative objective

$$J(u) := \alpha \|u\|^2. \quad (4.12)$$

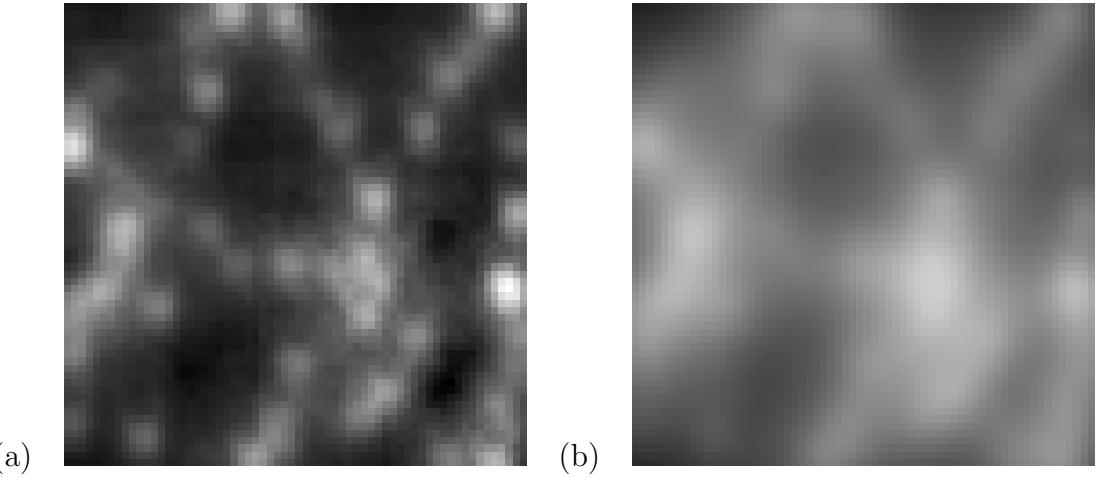
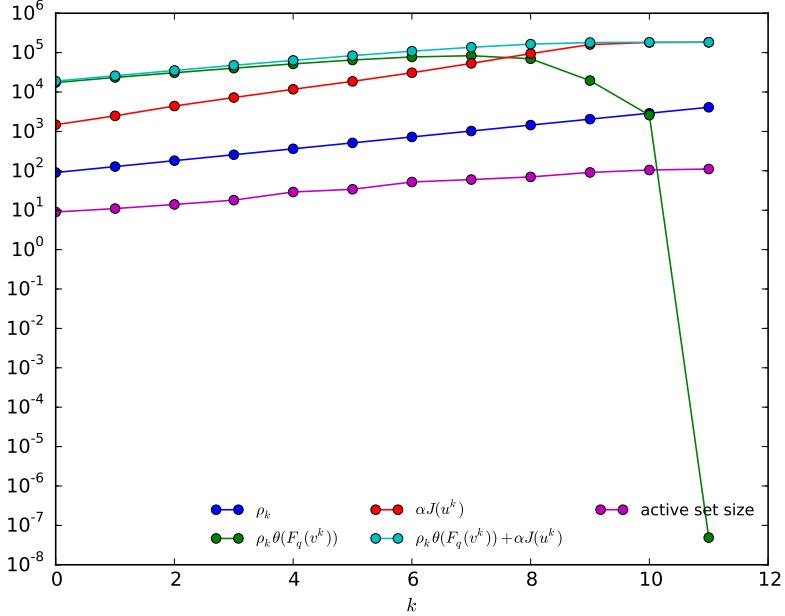


Figure 3: (a) Numerical reconstruction via Algorithm 4.1 from the imaging data shown in Figure 2 for $\rho = 4096$. (b) The reconstruction convolved with the measured PSF. At each resolution used for the reconstruction, the sum of the pixel values in (b) lie within a confidence interval of 3σ of those in Figure 2(b).

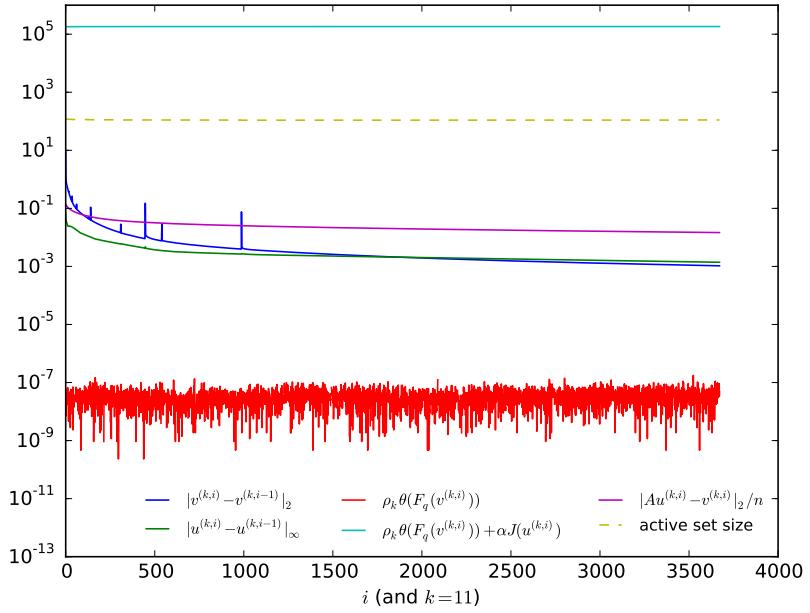
For the image size $n = 64 \times 64$ with the window system of squares of lengths 1 and 2, the number of windows is $M = 8065$. The constant α in (4.12) is, strictly speaking, redundant but was introduced as an additional means to balance the contributions of the individual terms to make the most of limited numerical accuracy (double precision). We chose $\alpha = 0.01$. The constant q was chosen so that the model solution would be no more than 3 standard deviations from the noisy data on each interval of each scale.

We emphasize that, since this is experimental data, there is no “truth” for comparison - the constraint, together with the error bounds on the numerical solution to the model solution provide statistical guarantees on the numerical reconstruction [26]. The numerical “image” generated from the reconstructed tubulin density, \bar{u} , is given by $\bar{v} = A\bar{u}$ and is shown in Figure 3(b); this figure is a denoised version of the measured data shown in Figure 2(b).

In Figure 4(a) a sample run of the algorithm shows a succession of outer iterations. The inner iteration is shown in Figure 4(b) with the value of $\rho_{11} = 4096$ for which the constraints are exactly satisfied (to within machine precision), indicating the correspondence of the computed solution of problem (\mathcal{P}_ρ) to a solution to the exact model problem (\mathcal{P}) . The eventual (starting from around iteration 1500) linear convergence of the algorithm can be seen in Figure 4(c). Under the assumption that the latter iterates are indeed in the region of local linear convergence, the observed convergence rate is $c = 0.9997$, which yields an a posteriori upper estimate of the pixelwise error of about $8.9062e^{-4}$, or 3 digits of accuracy at each pixel.



(a)



(b)

Figure 4: (a) Outer iterates k of Algorithm 4.1 showing solutions, constraint violation, the value of the regularizer, the objective value for the penalized problem (\mathcal{P}_ρ) and the active set size for successively larger values of the penalty parameter ρ . (b) Inner iterates of Algorithm 4.1 with $\rho_{11} = 4096$: step sizes, constraint violation, objective value and gap between the auxiliary, image-space variables $v^{(k,i)}$ and the primary, domain-space variables $u^{(k,i)}$.

5 Concluding remarks

We have focused our attention on the ADMM algorithm due partly to its prevalence in practice, and partly its amenability to our theoretical techniques. The parameter η in Algorithm 2.1 was left constant. How to choose this parameter in the context of minimization is a perplexing question and worthy of further study. Our theoretical framework can also be adapted to Krasnoselski-Mann relaxations of the Douglas–Rachford algorithm. Statements about this will appear in work underway studying more generally *averaged mappings*.

The statistical interpretation of the reconstruction in Figure 3(b) as described in [26, 27] opens the door to a quantitative approach to image processing, but this is only valid when one can estimate the distance of the numerical approximation to the exact solution to the underlying model optimization problem (\mathcal{P}) . Determining quantitative estimates for how close the numerical solution shown in Figure 3(a) is to an exact solution to problem (\mathcal{P}) under the assumption of exact evaluation of the associated prox operators has been the topic of our study.

What is needed and largely missing in the current treatment of algorithms in the literature is a complete error analysis accounting for accumulated errors at each stage of algorithms – due to finite precision or finite termination of iterative procedures – together with statements about how close one can get to the *solution* to a given optimization problem, as opposed to its *optimal value*, the latter having in general no necessary connection to the former. This is a monumental project that has not received as much attention in the literature as studies of complexity based upon function values. As we argued, the standard approach for handling inexact computation by assuming summable errors does not solve the problem, it just distributes it over infinitely many iterates. An alternative to this was suggested in [40, Section 6] and applied in [45] which allows a fixed error over all iterations without compromising local linear convergence. More work in this direction would narrow the gap between theory and practice.

Appendix

Duality of ADMM and the Douglas–Rachford Algorithm. Consider the sequence $(b^k, v^k)_{k \in \mathbb{N}}$ of the Douglas–Rachford iteration 2.8, for the case $B := \partial(J^* \circ (-A^T))$; $D := \partial H^*$. Recalling the two-step implementation (2.13), denote $\bar{p} := b^k - \eta v^k$ and $p' := q^{k+1}$. Then (2.13a) is the proximal step $p' = (I + \eta \partial(J^* \circ (-A^T)))^{-1} \bar{p}$ on the operator $B = \partial(J^* \circ (-A^T))$. If A has full column rank, by [22, Proposition 3.32(iv)], this step can be performed by

$$u^{k+1} = \arg \min_u \{J(u) + \langle \bar{p} + \eta v^k, Au \rangle + \frac{\eta}{2} \|Au - v^k\|^2\}; \quad (5.1)$$

$$p' = \bar{p} + \eta A u^{k+1}. \quad (5.2)$$

Indeed, since A has full rank, $J(u) + \langle \bar{p} + \eta v^k, Au \rangle + \frac{\eta}{2} \|Au - v^k\|^2$ is a proper strongly convex function of u and has a unique minimizer u^{k+1} . From the optimality condition

for (5.1),

$$0 \in \partial J(u^{k+1}) + A^T(\bar{p} + \eta Au^{k+1}) = \partial J(u^{k+1}) + A^T p'.$$

Hence, $(u^{k+1}, -A^T p') \in \text{gph } \partial J$ which implies $(-A^T p', u^{k+1}) \in \text{gph } \partial J^*$. This gives

$$\begin{aligned} &\Leftrightarrow (p', u^{k+1}) \in \text{gph } (\partial J^* \circ (-A^T)) \\ &\Leftrightarrow (p', -Au^{k+1}) \in \text{gph } (-A \circ \partial J^* \circ (-A^T)) \subseteq \text{gph } \partial (J^* \circ (-A^T)). \end{aligned}$$

Using (5.2),

$$\begin{aligned} (p', \frac{1}{\eta}(\bar{p} - p')) &\in \text{gph } \partial (J^* \circ (A^T)) \\ \Leftrightarrow p' &= (I + \eta \partial (J^* \circ (A^T)))^{-1} \bar{p}. \end{aligned}$$

Substituting $\bar{p} = b^k - \eta v^k$ in (5.1)-(5.2) yields

$$u^{k+1} = \arg \min_u \{J(u) + \langle b^k - \eta v^k + \eta v^k, Au \rangle + \frac{\eta}{2} \|Au - v^k\|^2\}; \quad (5.3)$$

$$q^{k+1} = b^k - \eta v^k + \eta Au^{k+1}. \quad (5.4)$$

Similarly, if we denote $\bar{p} := q^{k+1} + \eta v^k (= b^k + \eta Au^{k+1})$ and $p' := b^{k+1}$, (2.13b) is the proximal step $p' = (I + \eta \partial H^*)^{-1} \bar{p}$ on the operator $D = \partial H^*$ which can be performed via

$$\begin{aligned} v^{k+1} &= \arg \min_v \{H(v) - \langle \bar{p} - \eta Au^{k+1}, v \rangle + \frac{\eta}{2} \|Au^{k+1} - v\|^2\}; \\ p' &= \bar{p} - \eta v^{k+1}. \end{aligned}$$

Substituting $\bar{p} = b^k + \eta Au^{k+1}$,

$$v^{k+1} = \arg \min_v \{H(v) - \langle b^k + \eta Au^{k+1} - \eta Au^{k+1}, v \rangle + \frac{\eta}{2} \|Au^{k+1} - v\|^2\}; \quad (5.5)$$

$$b^{k+1} = b^k + \eta Au^{k+1} - \eta v^{k+1}. \quad (5.6)$$

Now, (5.3)-(5.4) and (5.5)-(5.6) together yield

$$\begin{aligned} u^{k+1} &= \arg \min_u \{J(u) + \langle b^k, Au \rangle + \frac{\eta}{2} \|Au - v^k\|^2\}; \\ v^{k+1} &= \arg \min_v \{H(v) - \langle b^k, v \rangle + \frac{\eta}{2} \|Au^{k+1} - v\|^2\}; \\ b^{k+1} &= b^k + \eta(Au^{k+1} - v^{k+1}). \end{aligned}$$

This is the ADMM algorithm (2.6a)-(2.6c) for the primal problem (\mathcal{P}_λ) . \square

Acknowledgments

We thank Jennifer Schubert of the Laser-Laboratorium Göttingen for providing us with the STED measurements shown in Fig. 2. Thanks also to Jalal Fadili for fruitful discussions and helpful comments during the preparation of this work.

References

- [1] F. J. ARAGÓN ARTACHO, J. M. BORWEIN, AND M. TAM, *Recent results on Douglas–Rachford methods*, Serdica Math. J., 39 (2013), pp. 313–330.
- [2] T. ASPELMEIER, A. EGNER and A. MUNK, *Modern statistical challenges in high-resolution fluorescence microscopy*, Annu. Rev. Stat. Appl., 2 (2015), pp. 163–202.
- [3] H. ATTTOUCH AND J. PEYPOUQUET, *The rate of convergence of Nesterov’s accelerated forward-backward method is actually $o(k^{-2})$* , arXiv:1510.08740 (Nov. 2015).
- [4] J.-P. AUBIN AND H. FRANKOWSKA, *Set-valued analysis*, Birkhäuser, Boston, 1990.
- [5] H. H. BAUSCHKE AND P. L. COMBETTES, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, CMS Books in Mathematics, Springer-Verlag, New York, 2011.
- [6] H. H. BAUSCHKE, P. L. COMBETTES, AND D. R. LUKE, *Finding best approximation pairs relative to two closed convex sets in Hilbert spaces*, J. Approx. Theory, 127 (2004), pp. 178–192.
- [7] H. H. BAUSCHKE, D. NOLL, AND H. M. PHAN, *Linear and strong convergence of algorithms involving averaged nonexpansive operators*, J. Math. Anal. Appl., 421 (2015), pp. 1–20.
- [8] H. H. BAUSCHKE, J. Y. BELLO CRUZ, T. T. A. NGHIA, H. M. PHAN, AND X. WANG, *The rate of linear convergence of the Douglas–Rachford algorithm for subspaces is the cosine of the Friedrichs angle*, J. Approx. Theory, 185 (2014), pp. 63–79.
- [9] H. H. BAUSCHKE AND D. NOLL, *On the local convergence of the Douglas–Rachford algorithm*, Arch. Math.(Basel), 102 (2014), pp. 589–600.
- [10] A. BECK AND M. TEBOULLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.
- [11] D. P. BERTSEKAS, *Constrained Optimization and Lagrange Multiplier Methods*, Academic Press, New York, 1982.
- [12] J. M. BORWEIN AND A. S. LEWIS, *Convex Analysis and Nonlinear Optimization: Theory and Examples*, Springer Verlag, New York, 2nd edition, 2006.
- [13] J. M. BORWEIN AND M. TAM, *The cyclic Douglas–Rachford method for inconsistent feasibility problems*, J. Nonlinear Convex Anal., 16 (2015), pp. 537–584.

- [14] D. BOLEY, *Local linear convergence of the alternating direction method of multipliers on quadratic or linear programs*. SIAM J. Optim., 23 (2013), pp. 2183–2207.
- [15] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Found. Trend Machine Learning, 3 (2011), pp. 1–122.
- [16] J. V. BURKE AND M. C. FERRIS, *Weak sharp minima in mathematical programming*, SIAM J. Control. Optim., 31 (1993), pp. 1340–1359.
- [17] J. V. BURKE, *An exact penalization viewpoint of constrained optimization*, SIAM J. Control Optim., 29 (1991), pp. 968–998.
- [18] F. H. CLARKE, *Optimization and Nonsmooth Analysis*, Wiley, New York, 1983. Republished as Vol. 5, Classics in Applied Mathematics, SIAM, 1990.
- [19] A. R. CONN, N. I. M. GOULD, AND PH. L. TOINT, *Trust Region Methods*, SIAM Publications, Philadelphia, 2000.
- [20] A. L. DONTCHEV AND R. T. ROCKAFELLAR, *Implicit Functions and Solution Mappings*, Springer-Verlag, New York, 2009.
- [21] J. DOUGLAS AND H. H. RACHFORD, *On the numerical solution of heat conduction problems in two or three space variables*, Trans. Amer. Math. Soc., 82 (1956), pp. 421–439.
- [22] J. ECKSTEIN, *Splitting Methods for Monotone Operators with Applications to Parallel Optimization*, PhD thesis, MIT, 1989.
- [23] J. ECKSTEIN AND D. P. BERTSEKAS, *On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators*, Math. Program., 55 (1992), pp. 293–318.
- [24] J. ECKSTEIN AND W. YAO, *Understanding the convergence of the alternating direction method of multipliers: theoretical and computational perspectives*, Pac. J. Optim., 11(2015), pp. 619–644.
- [25] R. FLETCHER, *An ℓ_1 penalty method for nonlinear constraints*, Numerical Optimization 1984, P. T. Boggs, R. H. Byrd, and R. B. Schnabel, eds., SIAM, Philadelphia, 1985, pp. 26–40.
- [26] K. FRICK, P. MARNITZ, AND A. MUNK, *Statistical multiresolution Dantzig estimation in imaging: Fundamental concepts and algorithmic framework*, Electron. J. Stat., 6 (2012), pp. 231–268.
- [27] K. FRICK, P. MARNITZ, AND A. MUNK, *Statistical multiresolution estimation for variational imaging: with an application in Poisson-biophotonics*, J. Math. Imaging and Vision, 46 (2013), pp. 370–387.

- [28] M. P. FRIEDLANDER AND P. TSENG, *Exact regularization of convex programs*, SIAM J. Optim., 18 (2007), pp. 1326–1350.
- [29] D. GABAY, *Augmented Lagrangian methods: Applications to the solution of boundary-value problems*, in Applications of the Method of Multipliers to Variational Inequalities, North-Holland, 1983, pp. 299–331.
- [30] P. GISELSSON, *Tight global linear convergence rate bounds for Douglas-Rachford splitting*, arXiv:1506.01556 (June, 2015).
- [31] P. GISELSSON, *Tight linear convergence rate bounds for Douglas-Rachford splitting and ADMM*, arXiv:1503.00887 (March, 2015).
- [32] R. GLOWINSKI AND A. MARROCO, *Sur l'approximation, par éléments finis d'ordre un, et la resolution, par penalisation-dualité, d'une classe de problèmes de dirichlet non linéaires*, Revue Francaise d'Automatique, Informatique et Recherche Opérationnelle, 9(R-2) (1975), pp. 41–76.
- [33] S.-P. HAN AND O. L. MANGASARIAN, *Exact penalty functions in nonlinear programming*, Math. Program., 17 (1979), pp. 251–269.
- [34] B. HE AND X. YUAN, *On the $O(1/n)$ convergence rate of the Douglas–Rachford alternating direction method*, SIAM J. Numer. Anal., 50 (2012), pp. 700–709.
- [35] S. W. HELL and J. WICHMANN, *Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy*, Optics Letters, 19 (1994), pp. 780–782.
- [36] R. HESSE AND D. R. LUKE, *Nonconvex notions of regularity and convergence of fundamental algorithms for feasibility problems*, SIAM J. Optim., 23 (2013), pp. 2397–2419.
- [37] R. HESSE, D. R. LUKE, AND P. NEUMANN, *Alternating projections and Douglas–Rachford for sparse affine feasibility*, IEEE Trans. Signal Process., 62 (2014), pp. 4868–4881.
- [38] T. A. KLAR, S. JAKOBS, M. DYBA, A. ENGER and S. W. HELL, *Fluorescence microscopy with diffraction resolution barrier broken by stimulated emission*, Proc. Natl. Acad. Sci. USA, 97 (2000), pp. 8206–8210.
- [39] D. KLATTE AND B. KUMMER, *Optimization methods and stability of inclusions in Banach spaces*, Math. Program., 117 (2009), pp. 305–330.
- [40] A. S. LEWIS, D. R. LUKE, AND J. MALICK, *Local linear convergence for alternating and averaged nonconvex projections*, Foundations of Computational Mathematics, 9 (2009), pp. 485–513.
- [41] A. S. LEWIS AND M. L. OVERTON, *Nonsmooth optimization via quasi-Newton methods*, Math. Program., 141 (2013), pp. 135–163.

- [42] G. LI AND T. K. PONG, *Douglas–Rachford splitting for nonconvex feasibility problems*, Math. Program., accepted.
- [43] J. LIANG, G. PEYRÉ, J. FADILI, AND D. R. LUKE, *Activity identification and local linear convergence of Douglas–Rachford under partial smoothness*, in Proceedings of the SSVM 2015, 2015.
- [44] P. L. LIONS AND B. MERCIER, *Splitting algorithms for the sum of two non-linear operators*, SIAM J. Numer. Anal., 16 (1979), pp. 964–979.
- [45] D. R. LUKE, *Local linear convergence of approximate projections onto regularized sets*, Nonlinear Anal., 75 (2012), pp. 1531–1546.
- [46] O. L. MANGASARIAN, *Sufficiency of exact penalty minimization*, SIAM J. Control Optim., 23 (1985), pp. 30–37.
- [47] J. J. MOREAU, *Proximité et dualité dans un espace Hilbertian*, Bull. Soc. math. Fr., 93 (1965), pp. 273–299.
- [48] R. NISHIHARA, L. LESSARD, B. RECHT, A. PACKARD, AND M. JORDAN, *A general analysis of the convergence of ADMM*, arXiv:1502.02009 (Feb. 2015).
- [49] J. NOCEDAL AND S. WRIGHT, *Numerical Optimization*, Springer Verlag, New York, 2000.
- [50] B. O’DONOGHUE, T. GOLDSTEIN, AND S. SETZER, *Fast alternating direction optimization methods*, SIAM J. Imaging Sci., 7 (2014), pp. 1588–1623.
- [51] N. PARikh AND S. BOYD, *Proximal Algorithms*, Found. Trends Optim., 1 (2014), pp. 123–231.
- [52] P. PATRINOS, L. STELLA, AND A. BEMPORAD, *Douglas–Rachford splitting: complexity estimates and accelerated variants*, in 53rd IEEE Conference on Decision and Control, 2014, 4234–4239.
- [53] J.-P. PENOT, *Calculus Without Derivatives*, Springer, New York, 2013.
- [54] H. PHAN, *Linear convergence of the Douglas–Rachford method for two closed sets*, Optimization, accepted.
- [55] R. T. ROCKAFELLAR, *Monotone operators and the proximal point algorithm*, SIAM J. Control Optim., 14 (1976), pp. 877–898.
- [56] R. T. ROCKAFELLAR AND R. J. WETS, *Variational Analysis*, Grundlehren der mathematischen Wissenschaften, Springer-Verlag, Berlin, 1998.
- [57] B. F. SVAITER, *On weak convergence of the Douglas–Rachford method*, SIAM J. Control. Optim., 49 (2011), pp. 280–287.